



WALLACE H. COULTER SCHOOL OF ENGINEERING
Technology Serving Humanity

MEMORANDUM

From: Bill Jemison
To: Dr. Daniel Tam, ONR
Date: 1/31/2013

Subject: Progress Report –
FY13 Progress Report (10/1/2013– 12/31/2013).

This document provides a progress report on the project “Advanced Digital Signal Processing for Hybrid Lidar” covering the period of 10/1/2013– 12/31/2013.

20150309474

Progress Summary

This report focuses on progress made towards developing a frequency domain reflectometry (FDR) ranging approach for hybrid lidar. The effectiveness of the technique is being explored through a combination of simulation using the Navy's Rangefinder simulation tool and experiments. Performance as a function of turbidity and integration time are explored.

A number of real-time DSP algorithms to support the FDR approach were developed using FPGA hardware.

Experimental and simulation results to date have shown accurate ranging performance at standoff ranges of over 7.2 and 6.3 attenuation lengths, respectively. The application of blind signal separation to the FDR algorithm is expected to enhance this performance to over 10 attenuation lengths.

Additional data collection at higher turbidities is planned in January 2014 to fully explore the performance limitation of FD and the application of blind signal separation.

Previous work was reported at the Oceans '13 conference.

Frequency Domain Reflectometry

In an attempt overcome the unambiguous range and range precision tradeoffs described previously, we have adapted a technique from the fiber optic community known as frequency-domain reflectometry (FDR). This technique was originally developed in the 1980s for the purpose of characterizing fiber lasers [4,5]. In the decades since, FDR has seen extensive use as an inexpensive method of approximating the location of faults in long fiber optic cables [6-9]. Utilizing modulation bandwidths of several gigahertz, this method has been used by the fiber optics community to unambiguously range over several kilometers of fiber optic cable with range resolutions on the order of 10 to 20 centimeters. Thus this method has been shown to simultaneously achieve high precision and high unambiguous ranging.

The key steps behind the FDR method will be briefly discussed. First, a stepped-frequency signal is transmitted into the channel. This signal reflects off objects in the channel and is collected by the receiver. The receiver measures the magnitude and phase of this return signal for all transmitted frequencies. This information is used to construct the frequency spectra for the current state of the channel, which encodes information about the distance to any objects currently in the receiver's field of view in the form of complex sinusoids. The inverse Fourier transform is taken to convert these complex sinusoids into sharp peaks in the time domain, indicating the time-of-flight required for the signal to reach each object in the channel. An automatic detection algorithm then searches the data for the sharpest peak, which is assumed to correspond to the object position. Finally, the time-of-flight information is converted into range data through knowledge of the speed of light in the medium.

A series of RangeFinder simulations were performed to assess performance of the FDR method, with as many parameters as possible set to match the actual equipment planned for use. These experiments aim

to use the 3.6 m long tank at Patuxent River Naval Air Station. In the planned experiments, an FDR configuration with a bandwidth of 500 MHz and 64 tones will be used, which correspond to an unambiguous range of 14.4 m and a range precision of 5.63 cm, respectively. The phase function of Maalox® antacid was used as an input to the simulation program. In addition, the effect of integration time on algorithm performance was simulated using a noise model [10]. Results for a harbor-like turbidity are shown below in Figures 2-4. First, the range return at the maximum object distance is shown in Figure 1, where the backscatter is seen to have an amplitude of approximately 32 dB while the object peak's amplitude is about 24 dB. It is easy for a human operator to examine this plot and determine that an object peak is present, but this is more challenging in an automated approach due to the lower magnitude of the object peak. As shown in Figure 2, the automatic detection algorithm is able to produce reasonable range calculations for the entire set of simulated positions. This simulation result indicates that the algorithm can detect object position out to at least 3.6 m in a harbor-like turbidity of $c=2$, which corresponds to 7.2 attenuation lengths. Finally, Figure 3 shows the effects of integration time on the range error, illustrating the mean range error and the calculated standard deviations for integration times of 1, 10, and 100 microseconds. Variation is reduced as integration time is increased, as the averaging of more data allows for a reduction in the effects of shot noise on the range calculation. The desired error-bounds for an application would ultimately determine the appropriate integration time. From the simulation results in Figure 3, a tolerance of ± 10 cm would correspond to a 1 microsecond integration time, while increasing the integration time to 100 microseconds allows for a tolerance of approximately ± 5 cm.

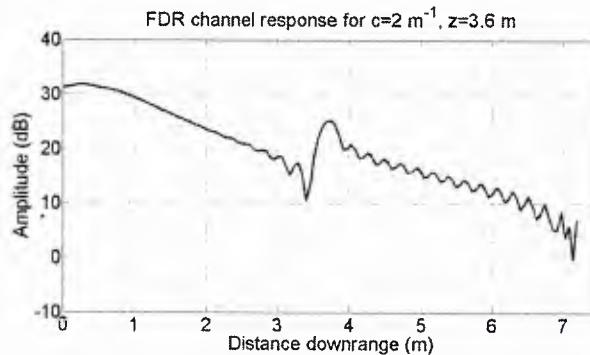


Figure 1. FDR channel response simulation for object at 3.6 m in harbor-like turbidity

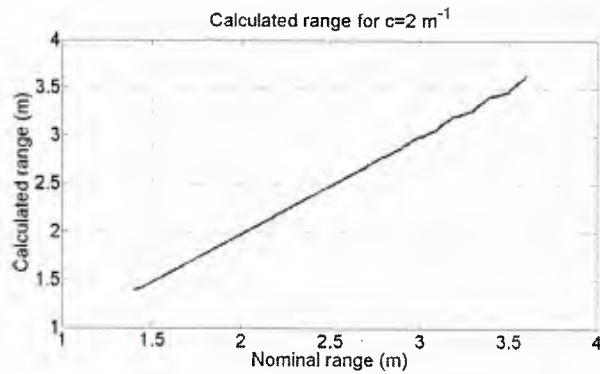


Figure 2. FDR range simulation for harbor-like turbidity

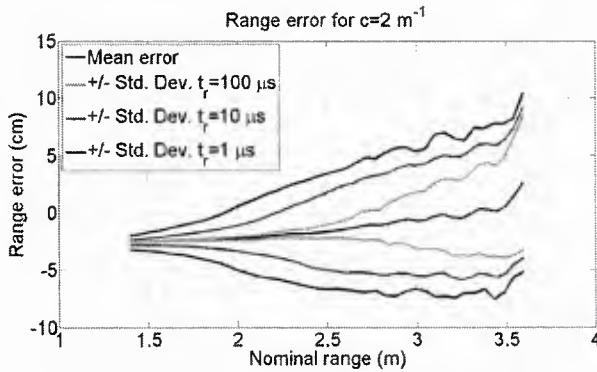


Figure 3. FDR range error simulation of integration time effects

A proof-of-concept experiment was performed in the 3.6 m tank in late summer, with additional experiments planned for late January 2014. The proof-of-concept experiment used the same FDR configuration as the simulations. The object position was fixed at the maximum distance while the turbidity was varied by dissolving Equate antacid into the tank. Results are shown below in Figure 4. The algorithm was able to calculate object positions out to 6.3 attenuation lengths, which corresponded to the maximum turbidity used in this set of experiments. This matches the simulation prediction that the algorithm can perform to at least 7.2 attenuation lengths. The deviation of the measured positions from the nominal line in Figure 4 at longer ranges also matches the simulated behavior, indicating that the accuracy of the range calculation begins to be affected by scattering at those longer ranges. Additional experiments will be performed to determine algorithm performance out to longer numbers of attenuation lengths in order to determine the maximum range at which the algorithm provides performance as well as to assess the expected performance degradation as the algorithm becomes backscatter-limited. These experiments will also be performed with various integration times to compare with the simulation results of Figure 3.

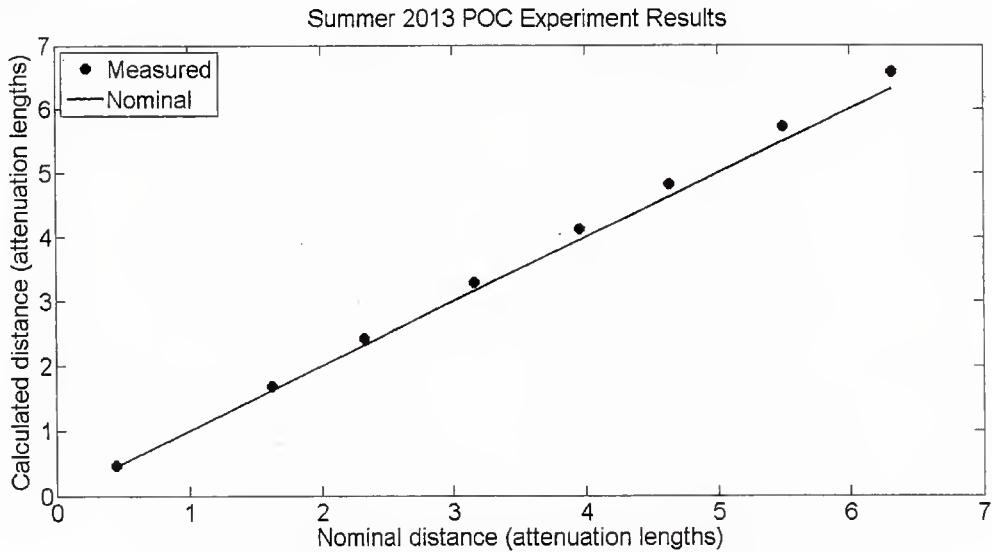


Figure 4. Results from summer 2013 proof-of-concept FDR experiments

Two practical concerns have come up regarding the FDR method since its introduction: integration time and processing requirements. First, the integration time issue is a result of the need for the transmitter to dwell on each of the frequencies in the stepped frequency sweep. The standard FDR processing approach cannot calculate a range result until the entire frequency sweep has completed. As an example, a FDR configuration using 64 tones would take 64 times as long to make a measurement as the single-tone CW approach. This concern is ultimately a technology issue rather than a problem with the FDR algorithm itself. The FDR technique was initially investigated in Spring 2013 using an inexpensive SDR-Kits network analyzer as the transmitter, for which the minimum dwell time was 1.33 milliseconds. For a configuration using 64 tones, this corresponds to 85.1 milliseconds per measurement or about 11 measurements per second. However, radar signal generators are capable of achieving much shorter dwell times, with the tradeoff being that these generators are more expensive. Radar signal generators are available that have dwell times in the tens or hundreds of microseconds, for example generators manufactured by Holzworth Instrumentation or FEI-Elcom Tech. As an example, for an FDR sweep using 64 tones, a radar signal generator with a minimum dwell time of one microsecond would produce over 15,000 range measurements per second. The minimum measurement rate in a given application informs the dwell time, which then allows for selection of an appropriate radar signal generator for use in the transmitter. As with many technology trends, it is expected that these radar signal generators will become less expensive over time.

The second practical concern relates to the processing requirements for calculating range with the FDR method, particularly related to the desire to implement the lidar ranging algorithms on real-time digital signal processing hardware. There are two main issues to consider with this in mind: processing time and resource utilization. With regard to these two metrics, the most expensive part of the FDR algorithm is the IFFT. Also, by moving towards a real-time version of the FDR method, additional logic is required to buffer results during the sweep, ensure that the magnitude and phase of each frequency is stored in the correct buffer location prior to the IFFT step, and ensure that all data have been searched before the peak detection algorithm outputs a range result. A particular challenge arises if frequency data are stored out of order, because this leads to an incorrect frequency spectrum input to the IFFT block and thus an incorrect range result will be calculated. These issues are essentially nonexistent when the FDR method is used as a post-processing tool, where a human operator knows what frequency the transmitter was using in every time index of every measurement. The resource utilization and frequency ordering problems motivated exploration of different ways of calculating a range result from the FDR method. One discrete object in the time-domain causes an impulse function to appear at the object location, which appears as a complex sinusoid in the frequency-domain by the properties of the Fourier Transform. This is why an IFFT is taken in the standard FDR approach; measurements were made in the frequency-domain and a processing algorithm (or human operator) expects to see impulse-like functions at the location of any objects in the time-domain. However, if the goal of the FDR sweep is to try to automatically detect one object position rather than assess the full channel response, the ranging problem can be reframed into a frequency demodulation problem. When an object is present in the channel in a target-limited scenario, the data measured in the FDR sweep are periodic, with a frequency that is a function of the object position. In scenarios where backscatter is present, lower frequencies in the FDR sweep begin to measure the volumetric center of the backscatter instead of the object position, but the higher frequencies resist scattering and maintain object information.

Two frequency demodulation approaches have been applied to previously simulated FDR data. In the first approach, the spacing between zero-crossings is used to estimate the object position. This approach is not particularly robust when the effects of shot noise are considered, as these might give rise to additional zero-crossings. A second approach based on the polar discriminator was also simulated, which is a technique that estimates the change in phase observed between two different measurements. This change in phase is directly proportional to the object position. Results for both of these approaches as well as the standard FDR method are shown below in Table 1 applied to simulated data with a turbidity of $c=1.6$. The standard IFFT-based FDR method is able to range to 4.7 m (7.52 a.l.) before becoming scatter-limited, while the zero-crossing method ranges to 5.3 m (8.48 a.l.) and the polar discriminator ranges to 5.6 m (8.96 a.l.). These simulation results indicate a possible advantage to performing processing in the frequency-domain where the measurements are made, as both of the frequency demodulation techniques can correctly range to farther object positions than the time-domain based processing of the standard IFFT-based FDR algorithm. These simulations do not include noise and are thus a representation of the ideal results; the zero-crossing method in particular is expected to have poorer performance in the presence of shot noise due to the expected presence of additional zero-crossings. The zero-crossing method has a mean range error that is more than double the mean error of the IFFT-based and polar discriminator approaches. Due to the IFFT and associated processing steps, the standard IFFT-based FDR method consumes significantly more resources and requires more clock cycles to calculate a result than the other two methods. In addition, the frequency demodulation techniques can begin to calculate range results after sweeping through as few as two frequencies, while the IFFT-based approach must wait for the sweep to complete so that the full frequency spectrum has been measured. This also means that the frequency demodulation methods do not need to maintain logic to buffer all sweep measurements in the correct order, because these methods can give a result for any arbitrary subset of the sweep measurements. As the frequencies in the sweep increase to be above the scatter-limited threshold, the range calculated will converge from an estimate of the backscatter-center to an estimate of the object position. The frequency demodulation approaches maintain the same unambiguous range as the IFFT technique, while potentially improving precision. This is because in the frequency demodulation approaches, the precision of a range calculation now depends on how accurately the magnitude and phase were measured, rather than on the sweep bandwidth. Table 1 indicates that there may be performance tradeoffs between the IFFT-based approach and a frequency demodulation approach, where some range may be lost or error may increase in order to use an algorithm that is less expensive in terms of both resource utilization and execution time. The data collected in the January experiments will be processed with each of these techniques to further assess their performance tradeoffs.

Table 1. Comparison of FDR processing methods

| Processing method | IFFT-based | Zero-crossing | Polar discriminator |
|---------------------------------|-------------------|-------------------|---------------------|
| Max. correct range (m) | 4.7 | 5.3 | 5.6 |
| Max. correct range (a.l.) | 7.52 | 8.48 | 8.96 |
| Mean range error (cm) | 3.55 | 8.11 | 2.68 |
| FPGA resource utilization* | 20.0% | 2.0% | 7.0% |
| FPGA clock latency | 5243 | 25 | 9 |
| Unambiguous range depends on... | frequency spacing | frequency spacing | frequency spacing |
| Range precision depends on... | sweep bandwidth | phase precision | phase precision |

* resource utilization specified as percentage of occupied slices on Xilinx Spartan-6 LX16 FPGA

Blind Signal Separation

For the FDR ranging approach, the statistical signal processing technique of blind signal separation (BSS) was adapted for backscatter reduction. In this technique, data are transformed into a statistical domain in which signals are separated based on their statistical properties [11]. This is analogous to using the Fourier transform to transform data into the frequency domain and separate signals based on their frequency content. Unlike the spatial filtering approach, BSS does not need to be adjusted for every modulation frequency, which made it a much more practical approach for backscatter suppression for the multiple frequencies required in the FDR method. A schematic of the BSS approach is shown below in Figure 5. In the top left, the frequency signal measured by FDR is shown, which contains both backscatter and target information. When this frequency signal is converted to range data, peaks for both the distributed backscatter and the target (correct location indicated with vertical green line) are obtained as shown in the top right. When BSS is applied to the frequency data, the scenario shown in the bottom left occurs, where the backscatter and target signals have been separated. By “zeroing out” the backscatter component, the range plot of the bottom right can be obtained, where the target still shows up in the correct position but the backscatter peak has been reduced by almost 10 dB. The BSS processing steps are critical in developing an automated target detection algorithm, such that the algorithm only detects a single peak instead of being confused by the backscatter return.

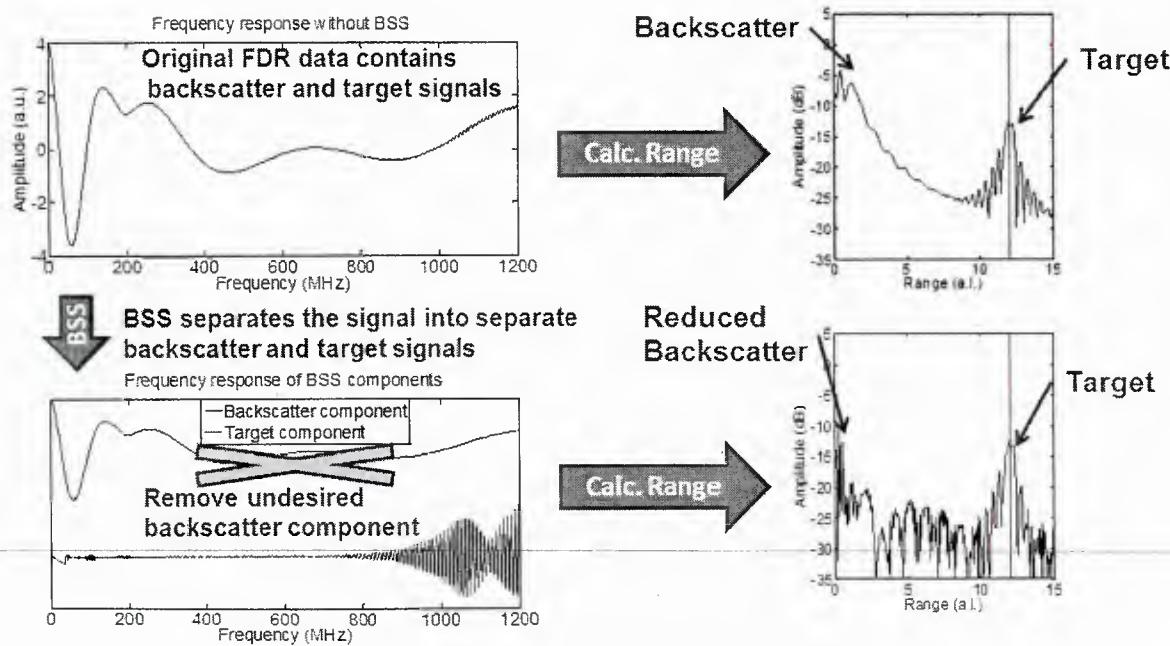


Figure 5. Schematic demonstrating approach for FDR with BSS

Blind signal separation will be applied to the data collected in the January experiments to assess the extent to which the technique can improve performance in a real system in the presence of noise. Simulations have been performed in RangeFinder over the range of attenuation lengths for the planned January experiments, with results shown below in Figure 6. The blue curve shows the range calculated by the standard FDR, which is able to correctly range out to 7.52 attenuation lengths (4.7 m). By applying BSS, an improvement of 10.4 attenuation lengths is observed, allowing the algorithm to range out to 17.92 attenuation lengths (11.2 m). This performance increase is achieved because the BSS technique has essentially cancelled out the backscatter return, allowing the target signal to be automatically detected. It is important to note that these simulations do not include receiver noise and thus represent a best-case scenario. The planned experiments will determine the performance improvement that can be expected in a real system.

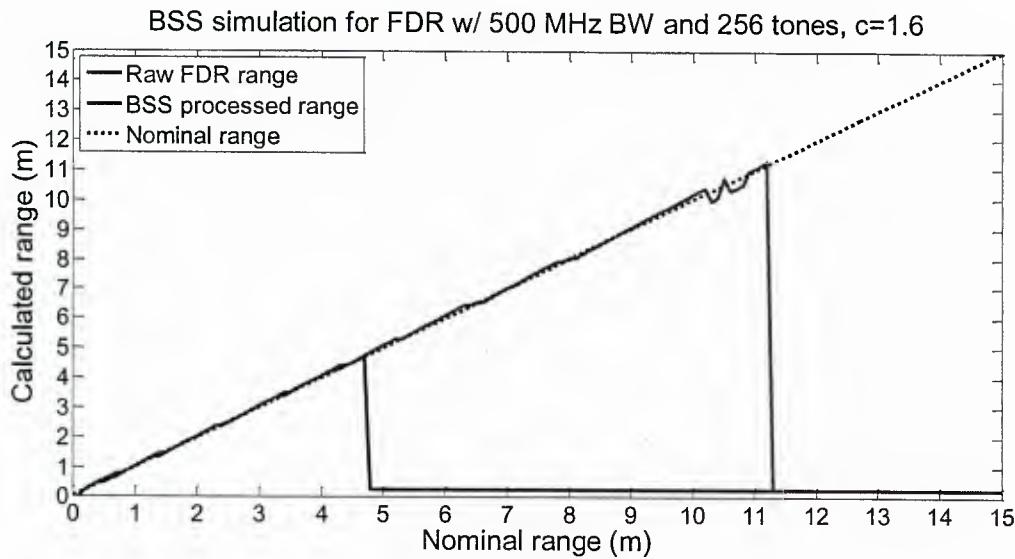


Figure 6. Blind signal separation simulation results

Deployment to Real-Time Systems

The second objective of this project is to deploy these ranging algorithms onto real-time digital signal processing (DSP) systems. Progress is summarized below in Table 2, where an X represents a completed task and the empty cells indicate remaining work, in some cases with an estimate of when that work will be completed. At this point the background information for each technique has been developed and small-scale experiments been performed for each technique. Large-scale tank experiments have been completed for the single-tone, dual-tone, and spatial filter techniques. The planned late January experiments will investigate the FDR, FDR/CW, and blind signal separation techniques in a large water tank. Real-time code has been designed for most of the techniques, with the code written and simulated in the Xilinx environment for the single-tone and dual-tone ranging algorithms. The next step will be to deploy these codes to an embedded system, at which point experiments will be performed to both validate the hardware system and to verify that the system produces meaningful results. Once the single-tone and dual-tone algorithms have been experimentally verified on real-time DSP hardware, the FDR algorithm will be developed and tested on the DSP hardware platform. The major focus will be on deploying the ranging algorithms to the DSP hardware platform, with the secondary goal being to also deploy the backscatter suppression algorithms. The methods will be compared in terms of both experimental performance and also resource utilization on the DSP hardware platform.

Table 2. Summary status of ranging algorithms

| | | Single-tone | Dual-tone | FDR | FDR/CW | Single-tone + spatial filter | Dual-tone + spatial filter | FDR + BSS |
|----------------------------|------------------------|-------------|-----------|------|--------|------------------------------|----------------------------|-----------|
| Background | Theory | X | X | X | X | X | X | X |
| | Simulation | X | X | X | X | X | X | X |
| Experimental Investigation | Proof-of-concept | X | X | X | X | X | X | X |
| | Small-scale (benchtop) | X | X | X | X | X | X | X |
| | Large-scale (tank) | X | X | Jan. | Jan. | X | X | Jan. |
| Real-Time Implementation | Code design | X | X | X | | X | X | |
| | Simulation | X | X | Mar. | | | | |
| | Implementation | Feb. | Feb. | Mar. | | | | |
| | Real-time experiment | Feb. | Feb. | Mar. | | | | |

Planned work and summary

Experiments are planned for late January to verify the FDR and BSS techniques under a variety of turbidities and practical system configurations which take into account shot noise limited performance. Following these experiments, the major emphasis will be on completing real-time implementations of the single-tone, dual-tone, and FDR ranging algorithms. This will include performing experiments to validate algorithm performance on these platforms.

The new FDR ranging approach has been shown to range out to 7.5 attenuation lengths in simulations that included a simulation of shot noise. Proof-of-concept experiments provide some evidence to support this ranging capability, as they showed that the algorithm could be used to range successfully out to at least 6.3 attenuation lengths. Processing improvements have been developed to reduce the calculation time and will be tested as part of the January experiments.

References

1. Mullen L.J., A.J.C. Vieira, and P.R. Herczfeld. "Application of RADAR technology to aerial LIDAR systems for enhancement of shallow underwater target detection," *IEEE Trans. Microw. Theory Techn.*, vol 43, pp. 2370-2377 (1995).

2. Pellen F., X. Intes, P. Olivard, Y. Guern, J. Cariou, and J. Lotrian. "Determine of sea-water cut-off frequency by backscattering transfer function measurement," *J. Phys. D: Appl. Phys.*, vol. 33, pp. 349-354 (2000).
3. Laux, A., L. Mullen, P. Perez, and E. Zege. "Underwater Laser Range Finder," *Proceedings of SPIE 2012, Ocean Sensing and Monitoring IV*, vol. 8372 (2012).
4. MacDonald, R.I. "Frequency domain optical reflectometer," *Applied Optics*, vol. 20, no. 10, pp. 1840-1844, May 1981.
5. Eickhoff, W. and R. Ulrich. "Optical frequency domain reflectometry in single-mode fiber," *Applied Physics Letters*, vol. 39, no. 9, pp. 693-695, Nov. 1981.
6. Healey, P. "Review of Long Wavelength Single-Mode Optical Fiber Reflectometry Techniques," *Journal of Lightwave Technology*, vol. LT-3, no. 4, Aug. 1985.
7. Ghafoori-Shiraz, H. and T. Okoshi. "Fault Location in Optical Fibers Using Optical Frequency Domain Reflectometry," *Journal of Lightwave Technology*, vol. LT-4, no. 3, Mar. 1986.
8. Barfuss, H. and E. Brinkmeyer. "Modified Optical Frequency Domain Reflectometry with High Spatial Resolution for Components of Integrated Optic Systems," *Journal of Lightwave Technology*, vol. 7, no. 1, pp. 3-10, Jan. 1989.
9. Furse, C., Y.C. Chung, C. Lo, P. Pendayala. "A critical comparison of reflectometry methods for location of wiring faults," *Smart Structures and Systems*, vol. 2, no. 1, pp. 25-46, 2006.
10. Illig, D., W.D. Jemison, L. Rumbaugh, R. Lee, A. Laux, and L. Mullen. "Enhanced hybrid lidar-radar ranging technique," *Proceedings of MTS/IEEE OCEANS 2013*, San Diego, California (2013).
11. A. Hyvärinen and E. Oja. "Independent component analysis: algorithms and applications," *Neural Networks*, vol. 13, no. 4/5, pp. 411-430 (2000).